

# Near-Linear Lower Bounds for Distributed Distance Computations, Even in Sparse Networks

Amir Abboud\*

Keren Censor-Hillel†

Seri Khoury†

May 18, 2016

## Abstract

We develop a new technique for constructing sparse graphs that allow us to prove near-linear lower bounds on the round complexity of computing distances in the CONGEST model. Specifically, we show an  $\tilde{\Omega}(n)$  lower bound for computing the diameter in sparse networks, which was previously known only for dense networks [Frishknecht et al., SODA 2012]. In fact, we can even modify our construction to obtain graphs with constant degree, using a simple but powerful degree-reduction technique which we define.

Moreover, our technique allows us to show  $\tilde{\Omega}(n)$  lower bounds for computing  $(\frac{3}{2} - \varepsilon)$ -approximations of the diameter or the radius, and for computing a  $(\frac{5}{3} - \varepsilon)$ -approximation of all eccentricities. For radius, we are unaware of any previous lower bounds. For diameter, these greatly improve upon previous lower bounds and are tight up to polylogarithmic factors [Frishknecht et al., SODA 2012], and for eccentricities the improvement is both in the lower bound and in the approximation factor [Holzer and Wattenhofer, PODC 2012].

Interestingly, our technique also allows showing an almost-linear lower bound for the verification of  $(\alpha, \beta)$ -spanners, for  $\alpha < \beta + 1$ .

---

\*Stanford University, Department of Computer Science, [abboud@cs.stanford.edu](mailto:abboud@cs.stanford.edu). Supported by Virginia Vassilevska Williams's NSF Grants CCF-1417238 and CCF-1514339, and BSF Grant BSF:2012338.

†Technion, Department of Computer Science, [{ckeren,serikhoury}@cs.technion.ac.il](mailto:{ckeren,serikhoury}@cs.technion.ac.il). Supported by ISF Individual Research Grant 1696/14.

# 1 Introduction

The diameter and radius are two basic graph parameters whose values play a vital role in many applications. In distributed computing, these parameters are even more fundamental, since they capture the minimal number of rounds needed in order to send a piece of information to all the nodes in a network. Hence, understanding the complexity of computing these parameters is central to distributed computing, and has been the focus of many studies in the CONGEST model of computation, where in every round each of  $n$  nodes may send messages of up to  $O(\log n)$  bits to each of its neighbors. Frischknecht et al. [21] showed that the diameter is surprisingly hard to compute:  $\tilde{\Omega}(n)$  rounds are needed even in networks with constant diameter.<sup>1</sup> This lower bound is nearly tight, due to an  $O(n)$  upper bound presented by [36] to compute all pairs shortest paths in a network. Naturally, approximate solutions are a desired relaxation, and were indeed addressed in several cornerstone studies [21, 24, 26, 30, 36], bringing us even closer to a satisfactory understanding of the time complexity of computing distances in distributed networks. However, several central questions remained elusive.

**Sparse Graphs.** The graphs constructed in [21] have  $\Theta(n^2)$  edges and constant diameter, and require any distributed protocol for computing their diameter to spend  $\tilde{\Omega}(n)$  rounds. Such a high lower bound makes one wonder if the diameter can be computed faster in networks that we expect to encounter in realistic applications. Almost all large networks of practical interest are very sparse [31], e.g. the Internet in 2012 had  $\approx 4$  billions nodes and  $\approx 128$  billion edges [32].

The only known lower bound for computing the diameter of a sparse network is obtained by a simple modification to the construction of [21] which yields a much weaker bound of  $\tilde{\Omega}(\sqrt{n})$ . This leaves hope that the  $\tilde{\Omega}(n)$  bound can be beaten significantly in sparse networks. Our first result is to rule out this possibility.

**Theorem 1.1.** *The number of rounds needed for any protocol to compute the diameter of a network on  $n$  nodes and  $O(n \log n)$  edges of constant diameter in the CONGEST model is  $\Omega(\frac{n}{\log^2 n})$ .*

We remark that, as in [21], our lower bound holds even for networks with constant diameter and even against randomized algorithms. Throughout the paper we say that a graph on  $n$  nodes is sparse if it has  $O(n \log n)$  edges. Due to simple transformations, e.g. by adding dummy nodes, all of our lower bounds will also hold for the more strict definition of sparse graphs as having  $O(n)$  edges, up to a loss of a log factor.

As explained next, the sparsity in our new lower bound construction allows us to extend the result in some interesting ways.

**Approximation Algorithms.** Another important question is whether we can bypass this near-linear barrier if we settle for knowing only an approximation to the diameter. An  $\alpha$ -approximation algorithm to the diameter returns a value  $\hat{D}$  such that  $D \leq \hat{D} \leq \alpha \cdot D$ , where  $D$  is the diameter of the network. From [21] we know that  $\tilde{\Omega}(\sqrt{n} + D)$  rounds are needed, even for computing a  $(\frac{3}{2} - \varepsilon)$ -approximation to the diameter, for any constant  $\varepsilon > 0$ .

Following this lower bound, almost-complementary upper bounds were under extensive research. It is known that a  $\frac{3}{2}$ -approximation can be computed in a sublinear number of rounds: Holzer and Wattenhofer [26] showed a  $O(n^{3/4} + D)$ -round algorithm and (independently) Peleg et al. [36] obtained a  $O(D\sqrt{n} \log n)$  bound, later these bounds were improved to  $O(\sqrt{n} \log n + D)$  by Lenzen and Peleg [30], and finally Holzer et al. [24] reduce the bound to  $O(\sqrt{n \log n} + D)$ . When  $D$  is

---

<sup>1</sup>The notations  $\tilde{\Omega}$  and  $\tilde{O}$  hide factors that are polylogarithmic in  $n$ .

small, these upper bounds are near-optimal in terms of the round complexity – but do they have the best possible approximation ratio that can be achieved within a sublinear number of rounds? That is, can we also obtain a  $(\frac{3}{2} - \varepsilon)$ -approximation in  $\tilde{O}(\sqrt{n} + D)$  rounds, to match the lower bound of [21]?

Progress towards answering this question was made by Holzer and Wattenhofer [26] who showed that any algorithm that needs to decide whether the diameter is 2 or 3 has to spend  $\tilde{\Omega}(n)$  rounds. However, as the authors point out, their lower bound is not robust and does not rule out the possibility of a  $(\frac{3}{2} - \varepsilon)$ -approximation when the diameter is larger than 2, or an algorithm that is allowed an additive +1 error besides a multiplicative  $(\frac{3}{2} - \varepsilon)$  error.

Perhaps the main difficulty in extending the lower bound constructions of [21] and [26] to resolve these gaps was that their original graphs are dense. A natural way to go from a lower bound construction against exact algorithms to a lower bound against approximations is to subdivide each edge into a path. When applied to dense graphs, this transformation blows up the number of nodes quadratically, resulting in a  $\tilde{\Omega}(\sqrt{n})$  lower bound [21]. Our new sparse construction technique allows us to tighten the bounds and negatively resolve the above question. In particular, we show a  $\tilde{\Omega}(n)$  lower bound for computing a  $(\frac{3}{2} - \varepsilon)$ -approximation to the diameter.

**Theorem 1.2.** *For all constant  $0 < \varepsilon < 1/2$ , the number of rounds needed for any protocol to compute a  $(3/2 - \varepsilon)$ -approximation to the diameter of a sparse network is  $\Omega(\frac{n}{\log^3 n})$ .*

**Radius.** In many scenarios we want one special node to be able to efficiently send information to all other nodes. In this case, we would like this node to be the one that is closest to every other node, i.e. the *center* of the graph. The *radius* of the graph is the largest distance from the center, and it captures the number of rounds needed for the center node to transfer a message to another node in the network. While radius and diameter are closely related, the previous lower bounds for diameter do not transfer to radius and it was conceivable that the radius of the graph could be computed much faster. Obtaining a non-trivial lower bound for radius has been stated as an open problem in [26]. A third advantage of our technique is that it extends to computing the radius, for which we show that the same strong near-linear barriers above hold.

**Theorem 1.3.** *For all constant  $0 < \varepsilon < 1/2$ , the number of rounds needed for any protocol to compute a  $(3/2 - \varepsilon)$ -approximation to the radius of a sparse network is  $\Omega(\frac{n}{\log^3 n})$ .*

**Eccentricity** The eccentricity of a node is the largest distance from it. Observe that the diameter is the largest eccentricity in the graph while the radius is the smallest. As pointed in [26], given a  $(\frac{3}{2} - \varepsilon)$ -approximation algorithm to all the eccentricities, we can achieve  $(\frac{3}{2} - \varepsilon)$ -approximation algorithm to the diameter by a simple flooding. This implies an  $\tilde{\Omega}(\sqrt{n} + D)$  lower bound for any  $(\frac{3}{2} - \varepsilon)$ -approximation algorithm for computing all the eccentricities. Our construction allows us to improve this result by showing that any algorithm for computing even a  $(\frac{5}{3} - \varepsilon)$ -approximation to all the eccentricities must spend  $\Omega(\frac{n}{\log^3(n)})$  rounds. This improves both in terms of the number of rounds, and in terms of the approximation factor, which we allow to be even larger. Interestingly, it implies that approximating all eccentricities is even harder than approximating just the largest or the smallest one.

**Theorem 1.4.** *For all constant  $0 < \varepsilon < 2/3$ , the number of rounds needed for any protocol to compute a  $(5/3 - \varepsilon)$ -approximation of all eccentricities of a sparse network is  $\Omega(\frac{n}{\log^3 n})$ .*

**Constant-Degree Graphs** For computing exact diameter and radius, we can modify the graph constructions according to a *degree-reduction technique* we define, such that the resulting graphs have a constant degree, and still allow us to obtain near-linear lower bounds. Roughly speaking, given a node  $v$ , we replace a subset of  $y$  edges of  $v$  by a binary tree to the respective neighbors (with additional internal nodes). This reduces the degree of  $v$  by  $y - 2$ . Repeatedly applying this procedure in a careful manner results in a graph of constant degree, for which we can still show our near-linear lower bounds (notice that distances change, as well as the number of nodes). We exemplify this technique by obtaining the following lower bound on computing the radius.

**Theorem 1.5.** *The number of rounds needed for any protocol to compute the radius of a sparse network of constant degree in the CONGEST model is  $\Omega(n/\log^3 n)$ .*

**Verification of Spanners** Finally, our technique allows us to obtain a lower bound for the verification of  $(\alpha, \beta)$ -spanners. An  $(\alpha, \beta)$ -spanner of a graph  $G$ , is a subgraph  $H$  in which for any two nodes  $u, v$  it holds that  $d_H(u, v) \leq \alpha d_G(u, v) + \beta$ . When spanners are sparse, i.e., when  $H$  does not have too many edges, they play a vital role in many application, such as routing, approximating distances, synchronization, and more. Hence, the construction of sparse spanners has been a central topic of many studies, both in centralized and sequential computing.

Here we address the problem of verifying that a given subgraph  $H$  is indeed an  $(\alpha, \beta)$ -spanner of  $G$ . At the end of the computation, each node outputs a bit indicating whether  $H$  is a spanner, with the requirement that if  $H$  is indeed a spanner with the required parameters then all nodes indicate this, and if it is not then at least one node indicates that it is not. We obtain the following.

**Theorem 1.6.** *Given an unweighted graph  $G = V, E$  and a subgraph  $H \subset E$  of  $G$ , the number of rounds needed for any protocol to decide whether  $H$  is an  $(\alpha, \beta)$ -spanner of  $G$  in the CONGEST model is  $\Omega(\frac{n}{(\alpha+\beta)\log^3 n})$ , for any  $\alpha < \beta + 1$ .*

Notice that for any reasonable value of  $\alpha, \beta = O(\text{poly } \log n)$ , the lower bound is near-linear. This is another evidence for a task for which verification is harder than computation in the CONGEST model, as initially brought into light in [17]. This is, for example, because  $(+2)$ -purely additive spanners with  $O(n^{3/2} \log n)$  edges can be constructed in  $O(\sqrt{n} \log n + D)$  rounds (this appears in [30], and can also be deduced from [26]), and additional various additive spanners can be constructed fast in CONGEST [13].

## 1.1 Techniques

**Communication Complexity and Distributed Computing.** A well-known technique to prove lower bounds in the CONGEST model is to use a reduction from communication complexity to distributed computing. Peleg and Rubinovich [37] apply a lower bound from communication complexity to show that the number of rounds needed for any distributed algorithm to construct a minimum spanning tree (MST) is  $\tilde{\Omega}(\sqrt{n} + D)$ . Many recent papers were inspired by this technique. In [19] Elkin extended the result of [37] to show that any distributed algorithm for constructing an  $\alpha$ -approximation to the MST must spend  $\tilde{\Omega}(\sqrt{\frac{n}{\alpha}})$  rounds. Das Sarma et al. [17] show that any distributed verification algorithm for many problems, such as connectivity,  $s-t$  cut and approximating MST requires  $\tilde{\Omega}(\sqrt{n} + D)$  rounds. Nanongkai et al. [34] showed an  $\Omega(\sqrt{\ell \cdot D} + D)$  lower bound for computing a random walk of length  $\ell$ . Similar reductions from communication complexity were adapted also in the CONGEST Clique Broadcast model [18, 25], where in each round each node can broadcast the same  $O(\log n)$ -bit message to all the nodes in the network.

Similar to the technique used in [17, 18, 21, 25, 26], our lower bounds are obtained by reductions from the Set-Disjointness problem in the two-party number-in-hand model of communication complexity [42]. Here, each of the players Alice and Bob receives a  $k$ -bit string,  $S_a$  and  $S_b$  respectively, and needs to decide whether the two strings are disjoint or not, i.e., whether there is some bit  $0 \leq i \leq k - 1$  such that  $S_a[i] = 1$  and  $S_b[i] = 1$ . A classical result [28, 38] is that in order to solve the Set-Disjointness problem, Alice and Bob must exchange  $\Omega(k)$  bits.

The high level idea for applying this lower bound in the *CONGEST* model, is as follows. We construct a graph in which the existence of some of the edges depends on the inputs of Alice and Bob, and we partition the graph between the two players, inducing a cut in it, which we will refer to as the “communication-cut”. The graph will have some property (e.g. diameter at least 4) if and only if the two strings of Alice and Bob are disjoint. The players can then simulate a distributed algorithm (e.g. for diameter), while exchanging only the bits that are sent by the algorithm on edges that belong to the communication-cut. If our cut has  $t$  edges, then Alice and Bob only exchange  $O(r \cdot t \cdot \log n)$  bits where  $r$  is an upper bound on the round complexity of the algorithm. Therefore, the lower bound on the communication complexity of Set-Disjointness implies a lower bound on the number of *rounds* required for any distributed algorithm (for diameter). Observe that the larger the communication-cut in the reduction, the smaller the lower bound for the distributed problem.

Having a *sparse graph with a small cut*, is what allows us to make this leap in the lower bounds. To achieve this, the key idea is to connect the nodes to a set of nodes that represent their binary value, and the only nodes on the cut are the nodes of the binary representation. We call this graph structure a *bit-gadget*, and it plays a central role in all of our graph constructions. This is inspired by graph constructions for different settings (e.g. [2], see additional discussion for sequential algorithms below).

## 1.2 Additional Related Work

There are many known upper [23, 29, 33] and lower [17, 20, 27, 37] bounds for approximate distance computation in *weighted* networks. For example, the weighted diameter of a network with underlying diameter  $D$  can be approximated to within  $(2 + o(1))$  in  $O(n^{1/2+o(1)} + D^{1+o(1)})$  rounds [23]. Moreover, such problems have also been considered in the *congested clique* model [12, 23, 33], where  $(1 + o(1))$ -approximate all pairs shortest paths can be computed in  $O(n^{0.158})$  rounds [12].

**Diameter and Radius in Sequential Algorithms.** Intuitively, the technical difficulty in extending the proof for diameter to work for radius as well is the difference in types between the two problems: the diameter asks for a pair of nodes that are far ( $\exists x \exists y$ ) while radius asks for a node that is close to everyone ( $\exists x \forall y$ ). Recent developments in the theory of (sequential) algorithms suggest that this type-mismatch could lead to fundamental differences between the two problems. Recall that classical *sequential* algorithms solve APSP in  $O(nm)$  time [16] and therefore both diameter and radius can be solved in quadratic  $O(n^2)$  time in sparse graphs.

Due to the lack of techniques for proving *unconditional* super-linear  $\omega(n)$  lower bounds on the runtime of sequential algorithms for any natural problem, a recent line of work seeks hardness results conditioned on certain plausible conjectures (a.k.a. “Hardness in P”). An interesting example of such result concerns the diameter: Roditty and Vassilevska W. [39] proved that if the diameter of sparse graphs can be computed in truly-subquadratic  $O(n^{2-\varepsilon})$  time, for any  $\varepsilon > 0$ , then the Strong Exponential Time Hypothesis (SETH) is false<sup>2</sup>, by reducing SAT to diameter. Since then, many

---

<sup>2</sup>SETH is a pessimistic version of the  $P \neq NP$  conjecture, which essentially says that CNF-SAT cannot be solved in  $(2 - \varepsilon)^n$  time. More formally, SETH is the assumption that there is no  $\varepsilon > 0$  such that for all  $k \geq 1$  we can solve  $k$ -SAT on  $n$  variables and  $m$  clauses in  $(2 - \varepsilon)^n \cdot \text{poly}(m)$  time.

other problems were shown to be “SETH-hard” (e.g. [1, 3, 5, 6, 8] to name a few) but whether a similar lower bound holds for radius is an open question [2, 4, 9, 10, 15, 39]. In fact, Carmosino et al. [11] show that there is a formal barrier for reducing SAT to radius<sup>3</sup>, and Abboud, Vassilevska W. and Wang [4] introduce a new conjecture to prove an  $n^{2-o(1)}$  lower bound for radius<sup>4</sup> (which has a similar  $\exists\forall$  type). Diameter and radius seem to behave differently also in the regime of dense and weighted graphs where the best known algorithms take roughly cubic  $O(n^3/2^{\sqrt{\log n}})$  time [14, 40] and it is known that radius can be solved in truly-subcubic  $O(n^{3-\varepsilon})$  time if and only if APSP can [2], but showing such a subcubic-equivalence between APSP and diameter is a big open question [2, 7, 41].

The framework and set-up in our *unconditional* lower bound proofs for distributed algorithms are very different from the ones in the works on *conditional* lower bounds for sequential algorithms discussed above. Still, some of our graph gadgets are inspired by the constructions in those proofs, e.g. [2, 4, 10, 15, 39]. Thus, it is quite surprising that our hardness proof for diameter transfers without much difficulty to a hardness proof for radius.

### 1.3 Model and basic definitions

We consider a synchronized network of  $n$  nodes represented by an undirected graph  $G = (V, E)$ . In each round, each node can send a different message of  $b$  bits to each of its neighbors. This model is known as the *CONGEST*( $b$ ) model, and as the *CONGEST* model when  $b = O(\log(n))$  [35].

The graph parameters that need to be computed are formally defined as follows.

**Definition 1.7.** (*Eccentricity, Diameter and Radius*) Let  $d(u, v)$  denote the length of the shortest path between the nodes  $u$  and  $v$ . The eccentricity  $e(u)$  of some node  $u$  is  $\max_{v \in V} d(u, v)$ . The Diameter (denoted by  $D$ ) is the maximum distance between any two nodes in the graph:  $D = \max_{u \in V} e(u)$ . The Radius (denoted by  $r$ ) is the maximum distance from some node to the “center” of the graph:  $r = \min_{u \in V} e(u)$ .

Finally, we define what we mean when we say that a graph is sparse.

**Definition 1.8.** (*sparse network*) A sparse network  $G = (V, E)$  is a network with  $n$  nodes and at most  $O(n \log(n))$  edges.

Recall, however, that all our results can be obtained for graphs that have a strictly linear number of  $\Theta(n)$  edges, at the cost of at most an additional  $O(\log n)$  factor in the lower bound.

**Roadmap.** Section 2 contains our lower bound for computing the exact or approximate diameter. In Sections 3, 4 and 6, we give our lower bounds for computing the exact or approximate radius, for computing eccentricities, and for verifying spanners, respectively. Section 5 uses our degree-reduction technique to show a graph construction with constant degree.

## 2 Computing the Diameter

In this section we present lower bounds on the number of rounds needed to compute the diameter *exactly* and *approximately* in sparse networks. First, in Section 2.1 we present a higher lower bound on the number of rounds needed for any algorithm to compute the exact diameter of a

<sup>3</sup>It would imply a new co-nondeterministic algorithm for SAT and refute the Nondeterministic-SETH, which is a strong version of  $\text{NP} \neq \text{CONP}$ .

<sup>4</sup>A truly-subquadratic algorithm for computing the radius of a sparse graph refutes the Hitting Set Conjecture: there is no  $\varepsilon > 0$  such that given two lists  $A, B$  of  $n$  subsets of a universe  $U$  of size  $\text{poly } \log n$  we can decide whether there is a set  $a \in A$  that intersects all sets  $b \in B$  in  $O(n^{2-\varepsilon})$  time.



Furthermore, we add four sets of nodes, which are our bit-gadget:  $F = \{f_j \mid 0 \leq j \leq \log(k) - 1\}$ ,  $T = \{t_j \mid 0 \leq j \leq \log(k) - 1\}$ ,  $F' = \{f'_j \mid 0 \leq j \leq \log(k) - 1\}$ ,  $T' = \{t'_j \mid 0 \leq j \leq \log(k) - 1\}$ , each of size  $\log(k)$ . We connect the sets  $F, T$  with  $F', T'$  by adding edges between  $f_i$  and  $t'_i$ , and between  $t_i$  and  $f'_i$ , for each  $0 \leq i \leq \log(k) - 1$ . To define the connections between the sets  $L, R$  and the sets  $F, T, F', T'$ , we add the following edges: For each  $\ell_i \in L$ , if  $i^j = 0$ , we connect  $\ell_i$  to  $f_j$ , otherwise, we connect  $\ell_i$  to  $t_j$ . Similarly, for each  $r_i \in R$ , if  $i^j = 0$  we connect  $r_i$  to  $f'_j$ , otherwise, we connect  $r_i$  to  $t'_j$ .

To complete the construction we add two additional nodes  $\{a, b\}$ . We connect  $a$  to all the nodes in  $F \cup T \cup \{\ell_k, \ell_{k+1}\}$ , and similarly, we connect  $b$  to all the nodes in  $F' \cup T' \cup \{r_k, r_{k+1}\}$ . We also add an edge between the nodes  $a$  and  $b$ .

**Claim 2.1.** *For every  $i, j \in [k - 1]$  it holds that  $d(\ell_i, r_j) = 3$  if  $i \neq j$ , and  $d(\ell_i, r_j) = 5$  otherwise.*

*Proof.* If  $i \neq j$ , there must be some bit  $h$ , such that  $i^h \neq j^h$ . Assume without loss of generality that  $i^h = 1$  and  $j^h = 0$ . Then,  $\ell_i$  is connected to  $t_h$  and  $r_j$  is connected to  $f'_h$ . Since  $t_h$  and  $f'_h$  are connected by an edge,  $d(\ell_i, r_j) = 3$ . For the second part of the claim, note that there are 4 options for any shortest path from  $\ell_i$  to  $r_i$ :

1. Through the nodes  $\ell_k, \ell_{k+1}, r_{k+1}, r_k$ .
2. Through some other node  $\ell_j \in L$  such that  $j \neq i$  in 2 steps, and then using the shortest path of length 3 between  $\ell_j$  and  $r_i$ .
3. Through some other node  $r_j \in R$  such that  $j \neq i$  in 3 steps, and then using the shortest path of length 2 between  $r_j$  and  $r_i$ .
4. Through at least one of the nodes  $a, b$ . Note that  $d(\ell_i, a) = 2$  and  $d(\ell_i, b) = 3$ . Similarly  $d(b, r_i) = 2$  and  $d(a, r_i) = 3$ .

Thus, any possible shortest path between  $\ell_i$  and  $r_i$  must have length 5. □

**Claim 2.2.** *For every  $u, v \in V \setminus (L \cup R)$  it holds that  $d(u, v) \leq 3$ .*

*Proof.* Let  $u, v \in V \setminus (L \cup R)$ . By definition,  $u$  is connected to one of the nodes in  $\{a, b\}$ . The same holds for  $v$  and since  $a$  and  $b$  are connected by an edge,  $d(u, v) \leq 3$ . □

**Corollary 2.3.** *For every  $u, v$  such that  $u \in (V_a \setminus L)$  or  $v \in (V_b \setminus R)$ , it holds that  $d(u, v) \leq 4$ .*

### 2.1.2 Reduction from Set-Disjointness

To prove Theorem 1.1, we show a reduction from the Set-Disjointness problem. Following the construction defined in the previous section, we define a partition  $(G_a = (V_a, E_a), G_b = (V_b, E_b))$ :

$$\begin{aligned} V_a &= L \cup F \cup T \cup \{\ell_k, \ell_{k+1}, a\}, E_a = \{(u, v) \mid u, v \in V_a \wedge (u, v) \in E\} \\ V_b &= R \cup F' \cup T' \cup \{r_k, r_{k+1}, b\}, E_b = \{(u, v) \mid u, v \in V_b \wedge (u, v) \in E\} \end{aligned}$$

The graph  $G_a$  is simulated by Alice and the graph  $G_b$  is simulated by Bob, i.e., in each round, all the messages that nodes in  $G_a$  send to nodes in  $G_b$  are sent by Alice to Bob. Bob forwards these messages to the corresponding nodes in  $G_b$ . All the messages from nodes in  $G_b$  to nodes in  $G_a$  are sent in the same manner. Each player receives an input string  $(S_a$  and  $S_b)$  of  $k$  bits. If the bit  $S_a[i] = 0$ , Alice adds an edge between the nodes  $\ell_i$  and  $\ell_{k+1}$ . Similarly, if  $S_b[i] = 0$ , Bob adds an edge between the nodes  $r_i$  and  $r_{k+1}$ .



**Observation 2.4.** For every  $u, v \in V_a$ , it holds that  $d(u, v) \leq 4$ . Similarly,  $d(u, v) \leq 4$  for every  $u, v \in V_b$ .

This is because  $d(u_a, \ell_{k+1}) \leq 2$  for any  $u_a \in V_a$ , and  $d(u_b, r_{k+1}) \leq 2$  for any node  $u_b \in V_b$ .

**Lemma 2.5.** The diameter of  $G$  is at least 5 iff the sets of Alice and Bob are not disjoint.

*Proof.* Consider the case in which the sets are disjoint i.e., for every  $0 \leq i \leq k-1$  either  $S_a[i] = 0$  or  $S_b[i] = 0$ . We show that for any  $u, v \in V$ , it holds that  $d(u, v) \leq 4$ . There are 3 cases:

1.  $u \in V_a$  and  $v \in V_a$ : By Observation 2.4,  $d(u, v) \leq 4$ .
2.  $u \in V_b$  and  $v \in V_b$ : By Observation 2.4,  $d(u, v) \leq 4$ .
3.  $u \in V_a$  and  $v \in V_b$ : Consider the case in which  $u \in (V_a \setminus L)$  or  $v \in (V_b \setminus R)$ . By Corollary 2.3  $d(u, v) \leq 4$ . Otherwise,  $u \in L$  and  $v \in R$ , and by Claim 2.1  $d(\ell_i, r_j) = 3$  for every  $i \neq j$ . Note that in case  $i = j$  it holds that  $d(\ell_i, r_i) \leq 4$ , since either  $\ell_i$  is connected to  $\ell_{k+1}$  or  $r_i$  is connected to  $r_{k+1}$ , and the shortest path between  $\ell_i$  and  $r_i$  is the one through the nodes  $\ell_{k+1}, r_{k+1}$ .

In case the two sets are not disjoint, there is some  $0 \leq i \leq k-1$  such that  $S_a[i] = 1$  and  $S_b[i] = 1$ . Note that in this case  $\ell_i$  is not connected directly by an edge to  $\ell_{k+1}$  and  $r_i$  is not connected directly by an edge to  $r_{k+1}$ . Therefore, there are only 4 options for any shortest path from  $\ell_i$  to  $r_i$ , which are stated in Claim 2.1, from which we get that in this case the diameter of  $G$  is at least 5.  $\square$

**Proof of Theorem 1.1** From Lemma 2.5, we get that any algorithm for computing the exact diameter of the graph  $G$  can be used to solve the Set-Disjointness problem. Note that since there are  $O(\log(k))$  edges in the cut  $(G_a, G_b)$ , in each round Alice and Bob exchange  $O(\log(k) \cdot \log(n))$  bits. Since  $k = \Omega(n)$  we deduce that any algorithm for computing the diameter of a network must spend  $\Omega(\frac{n}{\log^2(n)})$  rounds, and since  $|E| = O(n \log(n))$  this bound holds even for sparse networks.

## 2.2 $(\frac{3}{2} - \varepsilon)$ -approximation to the Diameter

In this Section we show how to modify our sparse construction to achieve a stronger lower bound on the number of rounds needed for any  $(\frac{3}{2} - \varepsilon)$ -approximation algorithm.

**Theorem 1.2** For all constant  $0 < \varepsilon < 1/2$ , the number of rounds needed for any protocol to compute a  $(3/2 - \varepsilon)$ -approximation to the diameter of a sparse network is  $\Omega(\frac{n}{\log^3 n})$ .

### 2.2.1 Graph Construction

The main idea to achieve this lower bound is to stretch our sparse construction by replacing some edges by paths of length  $P$ , an integer which will be chosen later. Actually, we only apply the following changes to the construction described in Section 2.1.1 (see also Figure 2 where  $P = 3$ ):

1. Remove the nodes  $a, b$  and their incident edges.
2. Replace all the edges incident to the nodes  $\ell_k, r_k$  by paths of length  $P$ .
3. Replace all the edges  $(u, v)$  such that  $u \in L$  and  $v \in (F \cup T)$  by paths of length  $P$ . Similarly, Replace all the edges  $(u, v)$  such that  $u \in R$  and  $v \in (F' \cup T')$  by paths of length  $P$ .




Figure 2: Graph construction,  $P = 3$  (diameter approximation). Some edges are omitted.

4. Add two additional sets  $L' = \{\ell'_i \mid 0 \leq i \leq k-1\}$ ,  $R' = \{r'_i \mid 0 \leq i \leq k-1\}$  each of size  $k$ . Connect each  $\ell'_i$  to  $\ell_i$ , and each  $r'_i$  to  $r_i$ , by a path of length  $P$ .

Furthermore, to simplify our proof, we connect each  $u \in (F \cup T)$  to  $\ell_{k+1}$  by a path of length  $P$ . Similarly, connect each  $u \in (F' \cup T')$  to  $r_{k+1}$  by a path of length  $P$ .

**Definition 2.6.** ( $Y(u, v)$ ) For each  $u, v \in V$  such that  $u$  and  $v$  are connected by a path of length  $P$ , denote by  $Y(u, v)$  the set of all nodes on the  $P$  path between  $u$  and  $v$  (without  $u$  and  $v$ ).

**Claim 2.7.** For every  $u, v \in V \setminus (L' \cup R' \cup_{i \in [k-1]} Y(\ell'_i, \ell_i) \cup_{i \in [k-1]} Y(r'_i, r_i))$  it holds that  $d(u, v)$  is at most  $4P + 1$ .

*Proof.* By the construction, the distance from  $u$  to one of the nodes in  $\{\ell_{k+1}, r_{k+1}\}$  is at most  $2P$ . The same holds for  $v$ . Thus,  $d(u, v) \leq 4P + 1$ .  $\square$

**Claim 2.8.** For every  $i, j \in [k-1]$  it holds that  $d(\ell'_i, r'_j) = 4P + 1$  if  $i \neq j$ , and  $d(\ell'_i, r'_j) = 6P + 1$  otherwise.

*Proof.* If  $i \neq j$ , there must be some bit  $h$ , such that  $i^h \neq j^h$ . Assume without loss of generality that  $i^h = 1$  and  $j^h = 0$ . This implies that  $\ell_i$  is connected to  $t_h$  by a path of length  $P$  and  $r_j$  is connected to  $f'_h$  by a path of length  $P$  as well. Since  $t_h$  and  $f'_h$  are connected by an edge,  $d(\ell_i, r_j) = 2P + 1$  and  $d(\ell'_i, r'_j) = 4P + 1$ . For the second part of the claim note that there are 3 options for any shortest path from  $\ell'_i$  to  $r'_i$ :

1. Through the node  $\ell_{k+1}$  in  $3P$  steps, and then using the shortest path of length  $3P + 1$  between  $\ell_{k+1}$  and  $r'_i$ .
2. Through some other node  $\ell_j \in L$ , such that  $j \neq i$  in  $3P$  steps, and then using the shortest path of length  $3P + 1$  between  $\ell_j$  and  $r'_i$ .

3. Through some other node  $r_j \in R$ , such that  $j \neq i$  in  $3P + 1$  steps, and then using the shortest path of length  $3P$  between  $r_j$  and  $r'_i$ .

Thus, for the second part of the claim, any possible shortest path between  $(\ell'_i, r'_i)$  must have length  $6P + 1$ .  $\square$

### 2.2.2 Reduction from Set-Disjointness

Following the above construction, we define a partition  $(G_a = (V_a, E_a), G_b = (V_b, E_b))$ :

$$\begin{aligned}
V_a &= \bigcup_{\substack{i \in [k-1] \\ j \in [\log(k)-1] \\ i^j=0}} Y(\ell_i, f_j) \bigcup_{\substack{i \in [k-1] \\ j \in [\log(k)-1] \\ i^j=1}} Y(\ell_i, t_j) \bigcup_{i \in [k-1]} Y(\ell'_i, \ell_i) \\
&\quad \bigcup_{i \in [k-1]} Y(\ell_i, \ell_k) \cup Y(\ell_k, \ell_{k+1}) \cup L' \cup L \cup F \cup T \cup \{\ell_k, \ell_{k+1}\} \\
E_a &= \{(u, v) | u, v \in V_a \wedge (u, v) \in E\} \\
V_b &= \bigcup_{\substack{i \in [k-1] \\ j \in [\log(k)-1] \\ i^j=0}} Y(r_i, f'_j) \bigcup_{\substack{i \in [k] \\ j \in [\log(k)-1] \\ i^j=1}} Y(r_i, t'_j) \bigcup_{i \in [k-1]} Y(r'_i, r_i) \\
&\quad \bigcup_{i \in [k-1]} Y(r_i, r_k) \cup Y(r_k, r_{k+1}) \cup R' \cup R \cup F' \cup T' \cup \{r_k, r_{k+1}\} \\
E_b &= \{(u, v) | u, v \in V_b \wedge (u, v) \in E\}
\end{aligned}$$

Each player receives an input string  $(S_a$  and  $S_b)$  of  $k$  bits. If  $S_a[i] = 0$ , Alice adds an edge between the nodes  $\ell_i$  and  $\ell_{k+1}$ . Similarly, if  $S_b[i] = 0$ , Bob adds an edge between the nodes  $r_i$  and  $r_{k+1}$ .

**Claim 2.9.** *Let  $0 \leq i \leq k - 1$  be such that  $S_a[i] = 0$  or  $S_b[i] = 0$ . Then the distance from the node  $\ell_i \in L$  to any node  $u \in (R \cup \{r_{k+1}\})$  is at most  $2P + 2$ .*

*Proof.* There are 3 cases:

1.  $u = r_{k+1}$ : Note that  $d(\ell_i, \ell_{k+1}) \leq 2P$  and  $d(\ell_i, r_{k+1}) \leq 2P + 1$ .
2.  $u = r_j \in R$  and  $j \neq i$ : By Claim 2.8  $d(\ell_i, r_j) = 2P + 1$ .
3.  $u = r_j \in R$  and  $j = i$ : Note that either  $\ell_i$  is connected to  $\ell_{k+1}$  directly by an edge or  $r_i$  is connected to  $r_{k+1}$  directly by an edge. Thus, one of the distances  $d(\ell_i, r_{k+1}), d(\ell_{k+1}, r_i)$  must be equal to 2, and both  $d(\ell_i, \ell_{k+1}), d(r_{k+1}, r_i)$  are at most equal to  $2P$ , thus,  $d(\ell_i, r_i) \leq 2P + 2$ .

$\square$

Note that any node in  $V_b$  is connected by a path of length at most  $P$  to some node in  $R \cup \{r_{k+1}\}$ , and any node in  $L'$  is connected by a path of length  $P$  to some node in  $L$ . Combining this with Claim 2.9 gives the following.

**Corollary 2.10.** *Let  $0 \leq i \leq k - 1$  be such that  $S_a[i] = 0$  or  $S_b[i] = 0$ . Then  $d(u, v_b) \leq 4P + 2$  for any  $u \in \{\ell'_i\} \cup Y(\ell'_i, \ell_i)$  and any  $v_b \in V_b$ . Symmetrically,  $d(u, v_a) \leq 4P + 2$  for any  $u \in \{r'_i\} \cup Y(r'_i, r_i)$  and any  $v_a \in V_a$ .*

**Lemma 2.11.** *The Diameter of  $G$  is  $6P + 1$  if the two sets of Alice and Bob are not disjoint, and  $4P + 2$  otherwise.*

*Proof.* Consider the case in which the two sets are not disjoint i.e., there is some  $0 \leq i \leq k-1$  such that  $S_a[i] = 1$  and  $S_b[i] = 1$ . Note that in this case  $\ell_i$  is not connected directly by an edge to  $\ell_{k+1}$  and  $r_i$  is not connected directly by an edge to  $r_{k+1}$ . Thus, there are only 3 options for any shortest path from  $\ell'_i$  to  $r'_i$ , both are stated in Claim 2.8, which implies that in this case  $d(\ell'_i, r'_i) = 6P + 1$ . Therefore, the diameter of  $G$  is at least  $6P + 1$ . Consider the case in which the sets are disjoint i.e., for each  $0 \leq i \leq k-1$ , either  $S_a[i] = 0$  or  $S_b[i] = 0$ , we need to prove that for every  $u, v \in V$ , it holds that  $d(u, v) \leq 4P + 2$ . There are 3 cases:

1.  $u \in V_a$  and  $v \in V_a$ : Note that  $d(u, \ell_k) \leq 2P$ , the same holds for  $v$ . Thus  $d(u, v) \leq 4P$ .
2.  $u \in V_b$  and  $v \in V_b$ : Note that  $d(u, r_k) \leq 2P$ , the same holds for  $v$ . Thus  $d(u, v) \leq 4P$ .
3.  $u \in V_a$  and  $v \in V_b$ : Consider the case in which  $u \in (V_a \setminus (L' \cup_{i \in [k-1]} Y(\ell'_i, \ell_i)))$  and  $v \in (V_b \setminus (R' \cup_{i \in [k-1]} Y(r'_i, r_i)))$ . By Claim 2.7  $d(u, v) \leq 4P + 1$ . Otherwise, by Corollary 2.10  $d(u, v) \leq 4P + 2$ .

□

**Proof of Theorem 1.2** To complete the proof we need to choose  $P$  such that  $(\frac{3}{2} - \varepsilon) \cdot (4P + 2) < (6P + 1)$ , this holds for any  $P > \frac{1}{2\varepsilon} - \frac{1}{2}$ . Note that  $k = \Omega(\frac{n}{\log(n)})$  for a constant  $\varepsilon$ . Thus, we deduce that any algorithm for computing  $(\frac{3}{2} - \varepsilon)$ -approximation to the diameter requires at least  $\Omega(\frac{n}{\log^3(n)})$  rounds. Furthermore, the number of nodes and the number of edges are both equal to  $\Theta(k \log(k) \cdot P)$ . Thus, this lower bound holds even for graphs with linear number of edges.

### 3 Computing the Radius

In this section we present lower bounds on the number of rounds needed to compute the radius *exactly* and *approximately* in sparse networks.

#### 3.1 Exact Radius

**Theorem 3.1.** *The number of rounds needed for any protocol to compute the radius of a sparse network of constant diameter in the CONGEST model is  $\Omega(n/\log^2 n)$ .*

As in the previous sections, first we describe a graph construction, and then apply a reduction from the Set-Disjointness problem.

##### 3.1.1 Graph construction

The graph construction for the radius is very similar to the one described in Section 2.1.1. We only apply the following changes to the construction described in Section 2.1.1 (see also Figure 3):<sup>6</sup>

1. Remove the nodes  $a, b$  and their incident edges.
2. For each  $0 \leq i \leq \log(k) - 1$  we add an edge between the nodes  $f_i \in F$  and  $t_i \in T$ . Similarly, we add an edge between  $f'_i \in F'$  and  $t'_i \in T'$ .
3. Add a small gadget which connects each  $\ell_i \in L$  to a new node  $x_3$  by a path of length 3  $(x_1, x_2, x_3)$ .

---

<sup>6</sup>Note that for the sake of simplicity, some of the edges are omitted from Figure 3.




Figure 3: Graph Construction (radius).

**Claim 3.2.** For every  $i, j \in [k-1]$  it holds that  $d(\ell_i, r_j) = 3$  if  $i \neq j$ , and  $d(\ell_i, r_j) = 4$  otherwise.

*Proof.* If  $i \neq j$ , there must be some bit  $h$ , such that  $i^h \neq j^h$ . Assume without loss of generality that  $i^h = 1$  and  $j^h = 0$ . Then,  $\ell_i$  is connected to  $t_h$  and  $r_j$  is connected to  $f'_h$ . Since  $t_h$  and  $f'_h$  are connected by an edge,  $d(\ell_i, r_j) = 3$ . For the second part of the claim, note that all the bits in the binary representation of  $i$  are the same as the bits in the binary representation of  $j$ . Thus, for any edge connecting  $\ell_i$  to some node  $f_h \in F$ , the corresponding edge from  $r_j$  is connected to the node  $f'_h \in F'$  and not to  $t'_h \in T'$ . Similarly, for any edge connecting  $\ell_i$  to some node  $t_h \in T$ , the corresponding edge from  $r_j$  is connected to the node  $t'_h \in T'$  and not to  $f'_h \in F'$ . Note that for every  $0 \leq h \leq \log(k) - 1$  the nodes  $f_h$  and  $f'_h$  are not connected directly by an edge and  $d(f_h, f'_h) = 2$  (the same holds for  $t_h$  and  $t'_h$ ). Therefore,  $d(\ell_i, r_j) = 4$  if  $i = j$ .  $\square$

### 3.1.2 Reduction from Set-Disjointness

The reduction is similar to the one used to prove Theorem 1.1. First we define a partition  $(G_a = (V_a, E_a), G_b = (V_b, E_b))$ :

$$\begin{aligned} V_a &= L \cup F \cup T \cup \{\ell_k, \ell_{k+1}, x_1, x_2, x_3\} \\ E_a &= \{(u, v) | u, v \in V_a \wedge (u, v) \in E\} \\ V_b &= R \cup F' \cup T' \cup \{r_k, r_{k+1}\} \\ E_b &= \{(u, v) | u, v \in V_b \wedge (u, v) \in E\} \end{aligned}$$

Each player receives an input string  $(S_a$  and  $S_b)$  of  $k$  bits. If  $S_a[i] = 1$ , Alice adds an edge between the nodes  $\ell_i$  and  $\ell_{k+1}$ . Similarly, if  $S_b[i] = 1$ , Bob adds an edge between the nodes  $r_i$  and

$r_{k+1}$ . Note that unlike the reduction from Section 2.1.2, here, each player adds some edge if the corresponding bit in its input string is 1, rather than 0.

**Observation 3.3.** *For every node  $u \in V \setminus L$  it holds that  $e(u) \geq 4$ .*

*Proof.* There are two cases:

1.  $u \in V \setminus (L \cup \{x_1, x_2, x_3\})$ : Note that  $d(u, x_3) \geq 4$ , since the only nodes that are connected to  $x_1$  are the nodes in  $L$ .
2.  $u \in \{x_1, x_2, x_3\}$ : Note that  $d(u, r_k) \geq 4$

□

**Lemma 3.4.** *The radius of  $G$  is 3 if and only if the two sets of Alice and Bob are not disjoint.*

*Proof.* Consider the case in which the two sets are disjoint. Note that for every  $0 \leq i \leq k-1$ , it holds that  $d(\ell_i, r_i) = 4$ , since either  $\ell_i$  is not connected directly by an edge to  $\ell_{k+1}$  or  $r_i$  is not connected directly by an edge to  $r_{k+1}$ . Combining this with Observation 3.3 and Claim 3.2 we get that for every  $u \in V$ ,  $e(u) \geq 4$ . Thus, the radius of  $G$  is at least 4 as well. In case the two sets are not disjoint, i.e., there is some  $0 \leq i \leq k-1$  such that  $S_a[i] = 1$  and  $S_b[i] = 1$ , we show that  $e(\ell_i) = 3$ : From Claim 3.2, for every  $0 \leq j \leq k-1$  such that  $i \neq j$ , it holds that  $d(\ell_i, r_j) = 3$ , and since  $S_a[i] = 1$  and  $S_b[i] = 1$ , it holds that  $d(\ell_i, r_i) = 3$  as well (the path through the nodes  $\ell_{k+1}$  and  $r_{k+1}$ ). It is straightforward to see that for every  $u \in V \setminus R$ , it holds that  $d(\ell_i, u) \leq 3$ . Therefore, the radius of  $G$  is  $\min_{u \in V} e(u)$  which is  $e(\ell_i) = 3$ . □

**Proof of Theorem 3.1** From Lemma 3.4, we get that any algorithm for computing the exact radius of the graph  $G$  solves the Set-Disjointness problem. Note that as the construction described in Section 2.1.1, there are  $O(\log(k))$  edges in the cut  $(G_a, G_b)$ . Thus, in each round Alice and Bob exchange  $O(\log(k) \cdot \log(n))$  bits. Since  $k = \Omega(n)$  we deduce that any algorithm for computing the exact radius of a network must spend  $\Omega(\frac{n}{\log^2(n)})$  rounds, and since  $|E| = O(n \log(n))$  this lower bound holds even for sparse networks.

### 3.2 $(\frac{3}{2} - \varepsilon)$ approximation to the Radius

We now show how to extend the construction described in Section 2.2.1 to achieve the same asymptotic lower bound for any  $(\frac{3}{2} - \varepsilon)$ -approximation algorithm to the radius as well.

**Theorem 1.3** *For all constant  $0 < \varepsilon < 1/2$ , the number of rounds needed for any protocol to compute a  $(3/2 - \varepsilon)$ -approximation to the radius of a sparse network is  $\Omega(\frac{n}{\log^3 n})$ .*

#### 3.2.1 Graph construction

Let  $G = (V, E)$  be the graph construction described in section 2.2.1. We describe the construction for this section in two steps. In the first step, we apply the following changes to  $G$ :

1. Replace the path of length  $P$  connecting  $\ell_k$  to  $\ell_{k+1}$  by a path of length  $2P$ . Similarly, replace the path of length  $P$  connecting  $r_k$  to  $r_{k+1}$  by a path of length  $2P$ .
2. Remove all the paths of length  $P$  connecting  $\ell_{k+1}$  with some node  $u \in (F \cup T)$ . Similarly, remove all the paths of length  $P$  connecting  $r_{k+1}$  with some node  $u' \in (F' \cup T')$ .




Figure 4: Graph Construction (radius approximation). Each edge between two sets of nodes represents all the edges between the corresponding sets.

Let  $V'$  denote the set of vertices  $V \setminus (L' \cup_{i \in [k-1]} Y(\ell'_i, \ell_i))$ , and let  $G'$  be the graph induced by the vertices in  $V'$ . Now we can describe the second step, in which we extend the construction in the following manner: we have two instances of  $G'$ , denoted by  $G'_1 = (V'_1, E'_1)$  and  $G'_2 = (V'_2, E'_2)$ . Each node in  $G'_1 \cup G'_2$  has an additional label, indicating whether it belongs to  $G'_1$  or  $G'_2$ . For example, the node  $\ell_i$  in  $G'_1$  is denoted by  $\ell_{(i,1)}$ , and the node  $r_i$  in  $G'_2$  is denoted by  $r_{(i,2)}$ . To complete the construction, for each  $0 \leq i \leq k-1$  we add two paths, each of length  $P$ , connecting  $\ell'_i$  to the nodes  $\ell_{(i,1)}, \ell_{(i,2)}$ .

### 3.2.2 Reduction from Set-disjointness

The reduction is similar to the one used to prove Theorem 1.2. First we define a partition  $(G_a = (V_a, E_a), G_b = (V_b, E_b))$ , let  $V'_a$  be the set of vertices defined by:

$$V'_a = L' \cup L_1 \cup L_2 \cup F_1 \cup F_2 \cup T_1 \cup T_2 \cup \{\ell_{(k,1)}, \ell_{(k+1,1)}, \ell_{(k,2)}, \ell_{(k+1,2)}\}$$

Similarly, let  $V'_b$  be the set of vertices defined by:

$$V'_b = R_1 \cup R_2 \cup R'_1 \cup R'_2 \cup F'_1 \cup F'_2 \cup T'_1 \cup T'_2 \cup \{r_{(k,1)}, r_{(k+1,1)}, r_{(k,2)}, r_{(k+1,2)}\}$$

Now we define the sets  $V_a, E_a, V_b, E_b$ :

$$\begin{aligned}
V_a &= V'_a \bigcup_{u,v \in V'_a} Y(u,v) \\
E_a &= \{(u,v) | u,v \in V_a \wedge (u,v) \in E\} \\
V_b &= V'_b \bigcup_{u,v \in V'_b} Y(u,v) \\
E_b &= \{(u,v) | u,v \in V_b \wedge (u,v) \in E\}
\end{aligned}$$

Each player receives an input string  $(S_a \text{ and } S_b)$  of  $k$  bits. If  $S_a[i] = 1$ , Alice adds two paths, each of length  $P$ , the first connects the nodes  $\ell_{(i,1)}$  and  $\ell_{(k+1,1)}$ , and the second connects the nodes  $\ell_{(i,2)}$  and  $\ell_{(k+1,2)}$ . Similarly, if  $S_b[i] = 1$ , Bob adds two paths of length  $P$ , the first connects the nodes  $r_{(i,1)}$ , and  $r_{(k+1,1)}$  and the second connects the nodes  $r_{(i,2)}$  and  $r_{(k+1,2)}$ .

**Observation 3.5.** *The node with the minimum eccentricity is in  $L'$ .*

This is because for any node  $u \in V'_1 \bigcup_{i \in [k-1]} Y(\ell'_i, \ell_{(i,1)})$  it holds that the node with  $\max_{v \in V} d(u,v)$  must be in  $V'_2$ . Similarly, for any node  $u \in V'_2 \bigcup_{i \in [k-1]} Y(\ell'_i, \ell_{(i,2)})$  it holds that the node with  $\max_{v \in V} d(u,v)$  must be in  $V'_1$ . Thus, any node  $u \in V \setminus L'$  must visit some node in  $L'$  in order to reach the node with  $\max_{v \in V} d(u,v)$ .

**Observation 3.6.** *For any  $0 \leq i \leq k-1$  and any  $u \in V_a$ , it holds that  $d(\ell'_i, u) \leq 4P$ .*

This is because the distance from any  $u \in V_a$  to one of the nodes  $\{\ell_{(k,1)}, \ell_{(k,2)}\}$  is at most  $2P$ , and  $d(\ell'_i, \ell_{(k,1)}) = d(\ell'_i, \ell_{(k,2)}) = 2P$  for any  $0 \leq i \leq k-1$ .

**Claim 3.7.** *Let  $0 \leq i \leq k-1$  be such that  $S_a[i] = 0$  or  $S_b[i] = 0$ . It holds that  $d(\ell'_i, r'_{(i,1)}) \geq 6P+1$ . Similarly,  $d(\ell'_i, r'_{(i,2)}) \geq 6P+1$  as well.*

*Proof.* Note that either  $\ell_{(i,1)}$  is not connected to  $\ell_{(k+1,1)}$  by a path of length  $P$  or  $r_{(i,1)}$  is not connected to  $r_{(k+1,1)}$  by a path of length  $P$ . Thus, there are three options for any shortest path between  $\ell'_i$  and  $r'_{(i,1)}$ :

1. Through the nodes  $\ell_{(k+1,1)}, r_{(k+1,1)}$ : In this case, the shortest path must visit at least one of the nodes  $\ell_{(k,1)}, r_{(k,1)}$ . Assume without loss of generality that the shortest path visits only  $\ell_{(k,1)}$ . Thus, the distance between  $\ell'_i$  and  $r'_{(i,1)}$  can be written as the sum of two distances:  $d(\ell'_i, r'_{(i,1)}) = d(\ell'_i, \ell_{(k,1)}) + d(\ell_{(k,1)}, r'_{(i,1)})$ . Note that  $d(\ell_{(k,1)}, r'_{(i,1)}) \geq 4P+1$  (and equals  $4P+1$  if  $S_b[i] = 1$ ), and  $d(\ell'_i, \ell_{(k,1)}) = 2P$ . Thus, we have  $d(\ell'_i, r'_{(i,1)}) \geq 2P + 4P + 1 = 6P + 1$ .
2. Through some other node  $\ell_{(j,1)} \in L_1$  such that  $j \neq i$  in  $3P$  steps, and then using the shortest path of length  $3P+1$  between  $\ell_{(j,1)}$  and  $r'_{(i,1)}$ .
3. Through some other node  $r_{(j,1)} \in L_1$  such that  $j \neq i$  in  $3P+1$  steps, and then using the shortest path of length  $3P$  between  $r_{(j,1)}$  and  $r'_{(i,1)}$ .

Thus,  $d(\ell'_i, r'_{(i,1)}) = 6P+1$ , and similarly,  $d(\ell'_i, r'_{(i,2)}) = 6P+1$  as well.  $\square$

Combining Claim 3.7 with Observation 3.5, we conclude the following.

**Corollary 3.8.** *The radius of  $G \geq 6P+1$  if the two sets of Alice and Bob are disjoint.*



**Claim 3.9.** *Let  $0 \leq i \leq k-1$  be such that  $S_a[i] = 1$  and  $S_b[i] = 1$ . Then the distance from  $\ell'_i$  to any node in  $V_b$  is at most  $4P+1$ .*

*Proof.* Note that  $\ell_{(i,1)}$  is connected to  $\ell_{(k+1,1)}$  by a path of length  $P$  and  $r_{(i,1)}$  is connected to  $r_{(k+1,1)}$  by a path of length  $P$ . Similarly,  $\ell_{(i,2)}$  is connected to  $\ell_{(k+1,2)}$  by a path of length  $P$  and  $r_{(i,2)}$  is connected to  $r_{(k+1,2)}$  by a path of length  $P$ . Thus, the distance from  $\ell'_i$  to each of the nodes  $r_{(k+1,1)}, r_{(k+1,2)}$  is  $2P+1$ , and  $d(\ell'_i, v) \leq 4P+1$  for any  $v \in (Y(r_{(k,1)}, r_{(k+1,1)}) \cup Y(r_{(k,2)}, r_{(k+1,2)}))$ . It remains to show that  $d(\ell'_i, v) \leq 4P+1$  for any  $v \in V_b \setminus (Y(r_{(k,1)}, r_{(k+1,1)}) \cup Y(r_{(k,2)}, r_{(k+1,2)}))$ . Note that  $\ell'_i$  can reach any node in  $R_1 \cup R_2$  in  $3P+1$  steps. Furthermore, by the definition of the construction, for any  $v \in V_b \setminus (Y(r_{(k,1)}, r_{(k+1,1)}) \cup Y(r_{(k,2)}, r_{(k+1,2)}))$  there is some node  $u \in (R_1 \cup R_2)$ , such that  $d(v, u) = P$ . Thus, for any  $v \in V_b \setminus (Y(r_{(k,1)}, r_{(k+1,1)}) \cup Y(r_{(k,2)}, r_{(k+1,2)}))$  it holds that  $d(\ell'_i, v) \leq 4P+1$ .  $\square$

Combining Observation 3.5, Observation 3.6, Corollary 3.8 and Claim 3.9, we conclude the following.

**Lemma 3.10.** *The Radius of  $G$  is  $4P+1$  if the two sets of Alice and Bob are not disjoint, and at least  $6P+1$  otherwise.*

**Proof of Theorem 1.3** To deduce Theorem 1.3 we need to choose  $P$  such that  $(\frac{3}{2} - \varepsilon) \cdot (4P+1) < (6P+1)$ , this holds for any  $P > \frac{1}{8\varepsilon} - \frac{1}{4}$ . Note that  $k = \Omega(\frac{n}{\log(n)})$  for a constant  $\varepsilon$ . Thus, we deduce that any algorithm for computing a  $(\frac{3}{2} - \varepsilon)$ -approximation to the radius requires at least  $\Omega(\frac{n}{\log^3(n)})$  rounds. Furthermore, the number of nodes and the number of edges are both equal to  $\Theta(k \log(k) \cdot P)$ . Thus, this lower bound holds even for graphs with linear number of edges.

### 3.3 Shaving an Extra Logarithmic Factor from the Denominator

In this section we show how to modify the construction described in Section 3.1.1 to achieve higher lower bounds for computing the radius *exactly* and *approximately*. The general idea is to expand the input strings of Alice and Bob while preserving the (asymptotic) size of the cut. We apply the following changes to the construction described in Section 3.1.1 (see also Figure 5):

1. Replace the node  $\ell_{k+1}$  by a set of nodes  $\{\ell_{k+1}^j \mid 0 \leq j \leq \log(k) - 1\}$  of size  $\log(k)$ . Similarly, replace the node  $r_{k+1}$  by a set of nodes  $\{r_{k+1}^j \mid 0 \leq j \leq \log(k) - 1\}$  of size  $\log(k)$ .
2. Connect each  $\ell_{k+1}^j$  with  $r_{k+1}^j$  by an edge.

For the graph partition, we add to  $V_a$  the nodes  $\{\ell_{k+1}^j \mid 0 \leq j \leq \log(k) - 1\}$  and we add to  $V_b$  the nodes  $\{r_{k+1}^j \mid 0 \leq j \leq \log(k) - 1\}$ . The input to the set-disjointness problem now differs from the previous construction, as follows. Each of the players Alice and Bob receives an input string  $(S_a, S_b)$  of size  $k \log(k)$ , each bit is represented by two indices  $i, j$  such that  $i \in \{0, \dots, k-1\}$  and  $j \in \{0, \dots, \log(k) - 1\}$ . If the bit  $S_a(i, j) = 1$  Alice adds an edge between the nodes  $\ell_i$  and  $\ell_{k+1}^j$ . Similarly, If the bit  $S_b(i, j) = 1$  Bob adds an edge between the nodes  $r_i$  and  $r_{k+1}^j$ . Note that if the two sets of Alice and Bob are disjoint, then for each  $0 \leq i \leq k-1$  it holds that  $d(\ell_i, r_i) = 4$ . Otherwise, there is some  $0 \leq i \leq k-1$  such that  $d(\ell_i, r_i) = 3$ . Thus, one can prove by case analysis, that the radius of  $G$  is 3 if and only if the two sets of Alice and Bob are not disjoint. Note that the size of the cut  $(G_a, G_b)$  remains  $\Theta(\log(k))$ ,  $k$  remains  $\Theta(n)$ , and the size of the strings is  $\Theta(n \log(n))$ . Thus, we achieve the following:




Figure 5: Graph Construction (radius),  $z = \log(k) - 1$ . Some edges are omitted, for clarity.

**Theorem 3.11.** *The number of rounds needed for any protocol to compute the radius of a sparse network of constant diameter in the CONGEST model is  $\Omega(n/\log n)$ .*

Similarly, we can apply the same idea to the construction described in Section 3.2.1 and achieve the following theorem:

**Theorem 3.12.** *The number of rounds needed for any protocol to compute a  $(3/2 - \varepsilon)$ -approximation to the radius of a sparse network is  $\Omega(n/\log^2 n)$ .*

## 4 Computing a $(\frac{5}{3} - \varepsilon)$ -approximation to the Eccentricity

In this section we show that any algorithm for computing a  $(\frac{5}{3} - \varepsilon)$ -approximation to all the eccentricities must spend  $\Omega(\frac{n}{\log^3(n)})$  rounds. Similar to the previous sections, we define a graph construction, and then apply a reduction from the set-disjointness problem.

### 4.1 Graph Construction

We simply apply the following changes on the construction described in Section 2.2.1 (see also Figure 6):

1. Remove all the nodes in  $L' \cup_{i \in [k-1]} Y(\ell'_i, \ell_i)$  and their incident edges.
2. Remove all the paths of length  $P$  connecting  $\ell_{k+1}$  with some node  $u \in (F \cup T)$ . Similarly, remove all the paths of length  $P$  connecting  $r_{k+1}$  with some node  $u' \in (F' \cup T')$ .



$\ell_i$  to any of the nodes in  $R$  is  $2P + 1$ , and to any of the nodes in  $R'$  is  $3P + 1$ . It is straightforward to see, by the definition of the construction, that the distance from  $\ell_i$  to any node in  $V \setminus R'$  is at most  $3P + 1$  as well.  $\square$

**Theorem 1.4** *For all constant  $0 < \varepsilon < 2/3$ , the number of rounds needed for any protocol to compute a  $(5/3 - \varepsilon)$ -approximation of all eccentricities of a sparse network is  $\Omega(\frac{n}{\log^3 n})$ .*

**Proof of Theorem 1.4** To complete the proof we need to choose  $P$  such that  $(\frac{5}{3} - \varepsilon) \cdot (3P + 1) < (5P + 1)$ , this holds for any  $P > \frac{2}{9\varepsilon} - \frac{1}{3}$ . Note that  $k = \Omega(\frac{n}{\log(n)})$  for a constant  $\varepsilon$ . Thus, we deduce that any algorithm for computing a  $(\frac{5}{3} - \varepsilon)$ -approximation to all the eccentricities requires at least  $\Omega(\frac{n}{\log^3(n)})$  rounds.

## 5 Networks with $\Delta = 5$

In this section we show how to modify our construction for radius computation to achieve almost the same lower bound for networks with  $\Delta = 5$ , where  $\Delta$  is the maximum degree of the network. Similar modification can be applied to the construction of the diameter as well. Consider the construction described in section 3.1. For each node  $v \in V$  such that  $v$  has a large degree, we replace a subset of the edges incident to  $v$  by a binary tree (see also Figure 7). For example, consider the node  $\ell_0$ . Instead of connecting  $\ell_0$  to each of its neighbors in  $F \cup T$  directly by edges, we add a binary tree of size  $O(\log(k))$ , whose root is  $\ell_0$  and whose leaves are the corresponding nodes in  $F \cup T$ . Note that there are  $O(k)$  nodes with degree  $O(\log(k))$  and  $O(\log(k))$  nodes with degree  $O(k)$  in the original graph.




Figure 7: Replacing  $y$  edges connected to some node  $v$  of degree  $y + z$  by a binary tree of size  $O(y)$ . Thus, the degree of  $v$  becomes  $z + 2$ .

**Theorem 1.5** *The number of rounds needed for any protocol to compute the radius of a sparse network of constant degree in the CONGEST model is  $\Omega(n/\log^3 n)$ .*

Now we describe the changes we apply to the construction described in Section 3.1.1 formally by 10 steps:

1. Remove the node  $x_3$  and its incident edge.
2. Replace the edges between  $x_1$  and  $L$  by a binary tree of size  $O(k)$  such that  $x_1$  is the root and the leaves are the nodes in  $L$ . Note that the height of this tree is exactly  $\log(k)$ .
3. Replace the edge  $(x_1, x_2)$  by a path of length  $\log(k) + 2\log\log(k) - 1$ . Denote the set of all the nodes on this path, including  $x_1$  and  $x_2$ , by  $P(x_1, x_2)$ .
4. Remove all the edges connecting some node in  $L \cup R$  with some node in the *bit-gadget*.
5. Connect each  $\ell_i$  in  $L$  to all the nodes that represent its binary value by a binary tree of size  $O(\log(k))$ , such that  $\ell_i$  is the root and the leaves are the corresponding nodes in  $F \cup T$ . Similarly, connect each  $r_i$  in  $R$  to all the nodes that represent its binary value by a binary tree of size  $O(\log(k))$ , such that  $r_i$  is the root and the leaves are the corresponding nodes in  $F' \cup T'$ . Note that the height of each such tree is exactly  $\log\log(k)$ .
6. After the previous step, each of the nodes in the *bit-gadget* is a leaf in  $k$  trees, thus, we replace each  $u$  in the *bit-gadget* by a binary tree of size  $O(k)$  such that the root is  $u$  and the leaves are its  $k$  parents in the  $k$  binary trees. Note that the height of each such tree is exactly  $\log(k)$ . Thus, after this step, the distance between  $\ell_0$  and  $f_0$ , for example, is  $\log\log(k) + \log(k) - 1$  and not  $\log\log(k) + \log(k)$ . This is because the parents of the leaves in the tree rooted at  $\ell_0$  are leaves in the trees rooted at the corresponding nodes in  $F \cup T$ .
7. Replace the edges in  $\{(\ell_i, \ell_k) \mid 0 \leq i \leq k-1\} \cup \{(r_i, r_k) \mid 0 \leq i \leq k-1\}$  by paths of length  $\log\log(k) - 1$ . Denote by  $P(\ell_k)$  the set of all the nodes on the paths connecting  $\ell_k$  to the nodes in  $L$ , including the nodes in  $L$  and the node  $\ell_k$ . Similarly, denote by  $P(r_k)$  the set of all the nodes on the paths connecting  $r_k$  to the nodes in  $R$ , including the nodes in  $R$  and the node  $r_k$ .
8. Replace the set of the first edges on every path that connects  $\ell_k$  to the nodes in  $L$  (i.e., the  $k$  edges that are incident to  $\ell_k$ ) by a binary tree of size  $O(k)$  and height exactly  $\log(k)$ . Similarly, replace set of the first edges on every path that connects  $r_k$  to the nodes in  $R$  (i.e., the  $k$  edges that are incident to  $r_k$ ) by a binary tree of size  $O(k)$  and height exactly  $\log(k)$ . Note that the edges  $(\ell_k, \ell_{k+1})$  and  $(r_k, r_{k+1})$  remain the same. Thus, after this step, the distance between  $\ell_0$  and  $\ell_k$ , for example, is  $\log\log(k) + \log(k) - 1$ , and thus, the distance between  $\ell_0$  and  $\ell_{k+1}$  is  $\log\log(k) + \log(k)$ .
9. Connect the node  $\ell_{k+1}$  to each of the nodes in  $L$  by a new path of length  $\log\log(k) - 1$ , for each  $\ell_i$ . Denote by  $q(\ell_i)$  the neighbor of  $\ell_i$  in the corresponding path of size  $\log\log(k) - 1$  connecting  $\ell_i$  to  $\ell_{k+1}$ . Similarly, connect the node  $r_{k+1}$  to each of the nodes in  $R$  by a new path of length  $\log\log(k) - 1$ , and for each  $r_i$ , denote by  $q(r_i)$  the neighbor of  $r_i$  in the corresponding path of size  $\log\log(k) - 1$  connecting  $r_i$  to  $r_{k+1}$ . Denote by  $P(\ell_{k+1})$  the set of all the nodes on the paths connecting  $\ell_{k+1}$  to the nodes in  $L$ , including the nodes in  $L$  and the node  $\ell_{k+1}$ . Similarly, Denote by  $P(r_{k+1})$  the set of all the nodes on the paths connecting  $r_{k+1}$  to the nodes in  $R$ , including the nodes in  $R$  and the node  $r_{k+1}$ .

10. Replace the set of the first edges on every path that connects  $\ell_{k+1}$  to the nodes in  $L$  (i.e., the  $k$  edges that are incident to  $\ell_{k+1}$ ) by a binary tree of size  $O(k)$  and height exactly  $\log(k)$ . Similarly, replace the first edges on every path that connects  $r_{k+1}$  to the nodes in  $R$  (i.e., the  $k$  edges that are incident to  $r_{k+1}$ ) by a binary tree of size  $O(k)$  and height exactly  $\log(k)$ .

We note that we keep the edges between the nodes  $f_i \in F$  and  $t_i \in T$  and the edges between the nodes  $f'_i \in F'$  and  $t'_i \in T'$  for each  $0 \leq i \leq \log(k) - 1$  as in Section 3.1. For each  $v \in V$  such that  $v$  is a root of some binary tree in  $G$ , denote by  $T(v)$  the set of all nodes in the binary tree rooted at  $v$  not including  $v$  itself.

**Claim 5.1.** *The maximum degree in the graph is 5.*

*Proof.* We show that for each  $v \in V$  the degree of  $v$  is at most 5. There are 7 cases:

1.  $v \in \{\ell_{k+1}, r_{k+1}\}$ : Note that  $\ell_{k+1}$  is a root of a binary tree, and other than nodes in  $T(\ell_{k+1})$ , it is connected only to the nodes  $\ell_k$  and  $r_{k+1}$ , thus the degree of  $\ell_{k+1}$  is 4. A similar argument holds for  $r_{k+1}$ .
2.  $v \in \{\ell_k, r_k\}$ : Note that  $\ell_k$  is a root of a binary tree, and other than nodes in  $T(\ell_k)$ , it is connected only to the node  $\ell_{k+1}$ , thus, the degree of  $\ell_k$  is 3. A similar argument holds for  $r_k$  as well.
3.  $v \in L \cup R$ : Note that for each  $\ell_i \in L$ , it holds that  $\ell_i$  is a root of a binary tree, and it is a leaf in the binary tree rooted at  $x_1$ . In addition, it is connected to  $q(\ell_i)$  and to another node connecting it to a path of length  $\log \log(k) - 1$ , which is connected to the binary tree rooted at  $\ell_k$ . Thus, its degree is 5. The same holds for each  $r_i \in R$ .
4.  $v \in P(x_1, x_2)$ : If  $v \in P(x_1, x_2) \setminus \{x_1\}$  then  $v$  is of degree 2. Otherwise, the degree of  $x_1$  is 3, since it is a root of a binary tree and it is connected to another node on the path  $P(x_1, x_2)$ .
5.  $v \in \text{bit-gadget}$ : Note that  $v$  is a root of a binary tree and it is connected to two additional nodes in the *bit-gadget*. Thus, its degree is 4.
6.  $v$  is an inner node in some binary tree: The degree of an inner node in a binary tree is at most 3.
7.  $v \in P(\ell_k) \cup P(r_k) \cup P(\ell_{k+1}) \cup P(r_{k+1}) \setminus (L \cup R)$ : All the nodes on these paths are of degree 2.

□

## 5.1 Reduction from Set Disjointness

Each player receives an input string ( $S_a$  and  $S_b$ ) of  $k$  bits. If  $S_a[i] = 0$ , Alice removes the edge connecting  $\ell_i$  to  $q(\ell_i)$ . Similarly, if  $S_b[i] = 0$ , Bob removes the edge connecting  $r_i$  to  $q(r_i)$ .

**Claim 5.2.** *For every  $0 \leq i, j \leq k - 1$  such that  $i \neq j$ , it holds that  $d(\ell_i, r_j) \leq 2 \log \log(k) + 2 \log(k) - 1$ .*

*Proof.* If  $i \neq j$ , there must be some bit  $h$ , such that  $i^h \neq j^h$ . Assume without loss of generality that  $i^h = 1$  and  $j^h = 0$ . Then, the distance between  $\ell_i$  and  $t_h$  is  $\log \log(k) + \log(k) - 1$ . Similarly, the distance between  $r_i$  and  $t'_h$  is  $\log \log(k) + \log(k) - 1$ . Note that  $f_h$  and  $t'_h$  are connected by an edge. Thus,  $d(\ell_i, r_j) \leq 2 \log \log(k) + 2 \log(k) + 2(-1) + 1 = 2 \log \log(k) + 2 \log(k) - 1$ . □

**Claim 5.3.** *If  $0 \leq i \leq k-1$  is such that  $S_a[i] = 0$  or  $S_b[i] = 0$ , then  $d(\ell_i, r_i) \geq 2 \log \log(k) + 2 \log(k)$ . Otherwise,  $d(\ell_i, r_i) = 2 \log \log(k) + 2 \log(k) - 1$ .*

*Proof.* If  $0 \leq i \leq k-1$  is such that  $S_a[i] = 0$  or  $S_b[i] = 0$ , then either  $\ell_i$  is not connect to  $q(\ell_i)$  directly by an edge, or  $r_i$  is not connect to  $q(\ell_i)$  directly by an edge. Thus, similar to the previous constructions there are two options for any shortest path between  $\ell_i$  and  $r_i$ . The first is through the *bit-gadget* and the second is through the nodes  $\ell_{k+1}, r_{k+1}$ , and both of them must have length at least  $2 \log \log(k) + 2 \log(k)$ . Otherwise, in case  $S_a[i] = 1$  and  $S_b[i] = 1$  it holds that  $d(\ell_i, \ell_{k+1}) = \log \log(k) + \log(k) - 1$ . Similarly,  $d(r_i, r_{k+1}) = \log \log(k) + \log(k) - 1$ . Thus,  $d(\ell_i, r_i) = 2 \log \log(k) + 2 \log(k) - 1$ .  $\square$

**Claim 5.4.** *If  $0 \leq i \leq k-1$  is such that  $S_a[i] = 1$  and  $S_b[i] = 1$ , then  $e(\ell_i) = 2 \log \log(k) + 2 \log(k) - 1$ .*

*Proof.* We show that for any  $v \in V$ , it holds that  $d(\ell_i, v) \leq 2 \log \log(k) + 2 \log(k) - 1$ . There are 2 cases:

1.  $v \in V_a$ : Here, there are 7 sub-cases:

- (a)  $v \in \{\ell_k, \ell_{k+1}\}$ : By the construction,  $d(\ell_i, \ell_k) = \log \log(k) + \log(k) - 1$ , and since  $S_a[i] = 1$ , it holds that  $d(\ell_i, \ell_{k+1}) = \log \log(k) + \log(k) - 1$  as well.
- (b)  $v \in L$ : In this case,  $\ell_i$  can use the node  $\ell_k$  in order to reach any node in  $L$  by a path of length  $2 \log \log(k) + 2 \log(k) - 2$ , since the distance between any  $\ell_i$  and  $\ell_k$  is  $\log \log(k) + \log(k) - 1$ .
- (c)  $v \in P(x_1, x_2) \cup T(x_1) \setminus L$ : Note that  $d(\ell_i, x_1) = \log(k)$ , and  $x_1$  can reach any node in  $T(x_1)$  in  $\log(k)$  steps, and any node in  $P(x_1, x_2)$  in  $\log(k) + 2 \log \log(k) - 1$  steps. Thus,  $d(\ell_i, v) \leq 2 \log \log(k) + 2 \log(k) - 1$ .
- (d)  $v \in F \cup T$ : Note that for each  $0 \leq j \leq \log(k) - 1$ , it holds that the distance from  $\ell_i$  to one of the nodes in  $\{f_j, t_j\}$  is  $\log \log(k) + \log(k) - 1$ , and since we keep the edges between the nodes  $f_j \in F$  and  $t_j \in T$  for each  $0 \leq j \leq \log(k) - 1$ , it holds that  $d(\ell_i, v) \leq \log(k) + \log \log(k)$ .
- (e)  $v \in \bigcup_{i \in [k-1]} T(\ell_i) \bigcup_{j \in [\log(k)-1]} T(f_j) \bigcup_{j \in [\log(k)-1]} T(t_j)$ : Note that  $v$  is at distance at most  $\log(k) + \log \log(k) - 2$  from some node in  $F \cup T$ . In this case  $d(\ell_i, v) \leq 2 \log(k) + 2 \log \log(k) - 2$  by case 1(d).
- (f)  $v \in T(\ell_k) \cup P(\ell_k) \setminus L$ : Note that  $v$  is at distance at most  $\log(k) + \log \log(k) - 1$  from  $\ell_k$ . In this case,  $d(\ell_i, v) \leq 2 \log(k) + 2 \log \log(k) - 2$  by case 1(a).
- (g)  $v \in T(\ell_{k+1}) \cup P(\ell_{k+1}) \setminus L$ : Note that  $v$  is at distance at most  $\log(k) + \log \log(k) - 1$  from  $\ell_{k+1}$ . In this case,  $d(\ell_i, v) \leq 2 \log(k) + 2 \log \log(k) - 2$  by case 1(a).

2.  $v \in V_b$ : Here, there are 6 cases:

- (a)  $v \in \{r_k, r_{k+1}\}$ : Since  $S_a[i] = 1$ ,  $d(\ell_i, r_{k+1}) = d(\ell_i, \ell_{k+1}) + d(\ell_{k+1}, r_{k+1}) = \log \log(k) + \log(k)$ , and  $d(\ell_i, r_k) = \log \log(k) + \log(k) + 1$ .
- (b)  $v \in R$ : By Claims 5.2 and 5.3, it holds that  $d(\ell_i, v) = 2 \log \log(k) + 2 \log(k) - 1$ .
- (c)  $v \in F' \cup T'$ : By case 1(d), it holds that  $d(\ell_i, v) \leq \log(k) + \log \log(k) + 1$ .
- (d)  $v \in \bigcup_{i \in [k-1]} T(r_i) \bigcup_{j \in [\log(k)-1]} T(f'_j) \bigcup_{j \in [\log(k)-1]} T(t'_j)$ : Note that  $v$  is at distance at most  $\log(k) + \log \log(k) - 2$  from some node in  $F' \cup T'$ . In this case  $d(\ell_i, v) \leq 2 \log(k) + 2 \log \log(k) - 1$  by case 2(c).

- (e)  $v \in T(\ell_k) \cup P(r_k) \setminus L$ : Note that  $v$  is at distance at most  $\log(k) + \log \log(k) - 1$  from  $r_k$ . In this case  $d(\ell_i, v) \leq 2 \log(k) + 2 \log \log(k) - 1$  by case 2(a).
- (f)  $v \in T(\ell_{k+1}) \cup P(r_{k+1}) \setminus L$ : Note that  $v$  is at distance at most  $\log(k) + \log \log(k) - 1$  from  $r_{k+1}$ . In this case  $d(\ell_i, v) \leq 2 \log(k) + 2 \log \log(k) - 1$  by case 2(a).

□

**Lemma 5.5.** *If the two sets of Alice and Bob are disjoint, then the radius of  $G$  is at least  $2 \log \log(k) + 2 \log(k)$ . Otherwise, there is some  $0 \leq i \leq k-1$  such that  $S_a[i] = 1$  and  $S_b[i] = 1$  and the radius of  $G$  is at most  $e(\ell_i) \leq 2 \log \log(k) + 2 \log(k) - 1$ .*

*Proof.* Consider the case in which the two sets are not disjoint. By Claim 5.3, for all the nodes in  $L$ , it holds that  $d(\ell_i, r_i) \geq 2 \log \log(k) + 2 \log(k)$ . Note that for all the nodes  $u \in V \setminus (L \cup T(x_1) \cup P(x_1, x_2))$ , it holds that  $d(u, x_2) \geq 2 \log \log(k) + 2 \log(k)$ , and for all the nodes  $u \in T(x_1) \cup P(x_1, x_2)$  it holds that  $d(u, r_i) \geq 2 \log \log(k) + 2 \log(k)$  for each  $0 \leq i \leq k-1$ . Thus, the radius of  $G$  is at least  $2 \log \log(k) + 2 \log(k)$  as well. The second part of the claim follows directly from Claim 5.4. □

**Proof of Theorem 1.5** As described before, we add  $O(k)$  binary trees of size  $O(\log(k))$ , and  $O(\log(k))$  binary trees of size  $O(k)$ , and  $O(k)$  paths of size  $O(\log \log(k))$ . Thus the total number of nodes added to the construction described in Section 3.1 is  $O(k \log(k))$ . Thus,  $k = \Omega(\frac{n}{\log(n)})$  and by Lemma 5.5 that any algorithm for computing the radius of  $G$  requires at least  $\Omega(\frac{n}{\log^3(n)})$  rounds, even in graphs with  $\Delta = 5$ .

**Remark 5.6.** *We remark that we believe that by a simple modification we can obtain the same result for graphs with maximum degree 3. This is done by replacing each node of degree 4 by 2 nodes, each of degree 3, and by replacing each node of degree 5 by 3 nodes, each with degree 3.*

## 6 Verification of Spanners

In this section we show that verifying an  $(\alpha, \beta)$ -spanner is also a hard task in the *CONGEST* model.

**Theorem 1.6** *Given an unweighted graph  $G = (V, E)$  and a subgraph  $H \subset E$  of  $G$ , the number of rounds needed for any protocol to decide whether  $H$  is an  $(\alpha, \beta)$ -spanner of  $G$  in the *CONGEST* model is  $\Omega(\frac{n}{(\alpha+\beta) \log^3 n})$ , for any  $\alpha < \beta + 1$ .*

In particular, this gives a near-linear lower bound when  $\alpha, \beta$  are at most polylogarithmic in  $n$ .

### 6.1 Graph Construction

Let  $P = \alpha x + \beta$ . For our proof for *unweighted* graphs we need  $x$  to be equal to 1. However, we keep the notation  $x$  to prove the same lower bound for *weighted* graphs as well. We apply the following changes to the construction described in Section 3 (see also Figure 8):<sup>7</sup>

1. Remove the nodes  $\{x_1, x_2, x_3\}$  and their incident edges.
2. Remove the nodes  $\{\ell_k, r_k\}$  and their incident edges.

---

<sup>7</sup>Some of the edges are omitted, for clarity.





*Proof.* According to Observation 6.1, it holds that  $d_H(\ell_i, r_i) = 2(\alpha x + \beta) + 1$ , while  $d_G(\ell_i, r_i) = 2x + 1$ . And since  $\alpha(2x + 1) + \beta < 2(\alpha x + \beta) + 1$  for any  $\alpha < \beta + 1$ , it holds that  $H$  is not an  $(\alpha, \beta)$  spanner of  $G$ .  $\square$

**Lemma 6.3.** *If the two sets of Alice and Bob are disjoint, then  $H$  is an  $(\alpha, \beta)$  spanner of  $G$  for any  $\alpha \geq 1$ .*

*Proof.* We show that  $d_H(u, v) \leq \alpha d_G(u, v) + \beta$  for any  $u, v \in V$ .

1.  $u = \ell_i \in L$  and  $v = r_j \in R$  such that  $i \neq j$ : Note that  $d_G(\ell_i, r_j) = \min(2x+1, \alpha x + \beta + 1)$ , while  $d_H(\ell_i, r_j) = \alpha x + \beta + 1$  by Observation 6.1. Thus, the problematic case is when  $d_G(\ell_i, r_j) = 2x + 1$ , for which  $\alpha(d_G(\ell_i, r_j)) + \beta \geq \alpha(2x + 1) + \beta = 2\alpha x + \alpha + \beta \geq \alpha x + \beta + 1 = d_H(\ell_i, r_j)$  for any  $\alpha \geq 1$ .
2.  $u = \ell_i \in L$  and  $v = r_j \in R$  such that  $i = j$ : Note that  $d_G(\ell_i, r_i) \geq x + \alpha x + \beta + 1$ , since either  $\ell_i$  is not connected to  $\ell_{k+1}$  by a path of length  $x$ , or  $r_i$  is not connected to  $r_{k+1}$  by a path of length  $x$ . While  $d_H(\ell_i, r_i) = 2(\alpha x + \beta) + 1$  by Observation 6.1. Note that  $\alpha(d_G(\ell_i, r_i)) + \beta \geq \alpha(x + \alpha x + \beta + 1) + \beta \geq \alpha x + \alpha^2 x + \alpha \beta + \alpha + \beta \geq \alpha x + \alpha x + \beta + 1 + \beta \geq 2(\alpha x + \beta) + 1 = d_H(\ell_i, r_i)$ , for any  $\alpha \geq 1$ .
3.  $u \in V_a$  and  $v \in V_a$ : Here, there are two cases:
  - (a) The shortest path in  $G$  does not visit the node  $\ell_{k+1}$ . In this case  $d_G(u, v) = d_H(u, v)$ .
  - (b) The shortest path in  $G$  visits the node  $\ell_{k+1}$ , thus, it visits some node  $\ell_i$ . In this case  $d_G(u, v)$  can be written as the sum of three distances,  $d_G(u, v) = d_G(u, \ell_i) + d_G(\ell_i, \ell_{k+1}) + d_G(\ell_{k+1}, v)$ , if none of the distances  $d_G(u, \ell_i), d_G(\ell_{k+1}, v)$  is using the disjointness edges, then the claim holds by the fact that  $\alpha(d_G(\ell_i, \ell_{k+1})) + \beta \geq d_H(\ell_i, \ell_{k+1})$ . Otherwise,  $d_G(u, v)$  can be written as the sum of three distances,  $d_G(u, v) = d_G(u, \ell_i) + d_G(\ell_i, \ell_j) + d_G(\ell_j, v)$ , note that  $d_G(\ell_i, \ell_j) \geq 2x$ , while  $d_H(\ell_i, \ell_j) = \alpha x + \beta$  (the path through the node  $\ell_{k+2}$ ), and the claim holds by the fact that  $\alpha 2x + \beta \geq \alpha x + \beta$ .
4.  $u \in V_b$  and  $v \in V_b$ : This case is symmetric to the previous one.
5.  $u \in V_a$  and  $v \in V_b$ : Here, there are four cases:
  - (a) The shortest path in  $G$  does not visit the nodes  $\ell_{k+1}, r_{k+1}$ . In this case  $d_G(u, v) = d_H(u, v)$ .
  - (b) The shortest path in  $G$  visits only the node  $\ell_{k+1}$  and does not visit  $r_{k+1}$ : This case is very similar to the case 3.b.
  - (c) The shortest path in  $G$  visits only the node  $r_{k+1}$  and does not visit  $\ell_{k+1}$ : This case is symmetric to the previous case.
  - (d) The shortest path visits the two nodes  $\ell_{k+1}, r_{k+1}$ , thus, it visits some two nodes  $\ell_i, r_j$ . In this case  $d_G(u, v)$  can be written as the sum of three distances,  $d_G(u, v) = d_G(u, \ell_i) + d_G(\ell_i, r_j) + d_G(r_j, v)$ , in case  $i \neq j$ , it holds that  $\alpha(d_G(u, v)) + \beta \geq d_H(u, v)$  by the case in which  $u = \ell_i \in L$  and  $v = r_j \in R$  such that  $i \neq j$  which was proved in the first case of this proof. Otherwise,  $\alpha(d_G(u, v)) + \beta \geq d_H(u, v)$  by the case in which  $u = \ell_i \in L$  and  $v = r_j \in R$  such that  $i = j$  which was proved in the second case of this proof.

$\square$

**Observation 6.4.** *Note that the number of edges  $E(G)$  is equal to  $\Theta(E(H))$ . However, it is straightforward to see that by adding a clique of size  $k$  to  $G$  and connecting it to some arbitrary node we can control the number of edges in  $G$ . We add  $k$  additional edges to  $H$  in order to span this clique, giving that the number of edges in  $G$  is  $\Theta(n^2)$ , while the number of edges in  $H$  is  $\Theta(n \log(n))$  (and if  $\beta > 2$  the number of edges in  $H$  is equal to  $\Theta(n)$ ). If we want  $E(H)$  to match the known bounds for the possible sparsity of an  $(\alpha, \beta)$ -spanner for general graphs, we simply add more edges of the clique to  $H$ .*

**Proof of Theorem 1.6** From the two lemmas 6.2 and 6.3 we deduce that  $H$  is an  $(\alpha, \beta)$ -spanner of  $G$  if and only if the two sets of Alice and Bob are disjoint. Note that for *unweighted* graphs we need  $x$  to be equal to 1, thus,  $n = O(k \cdot (\alpha + \beta) \log(n))$ , i.e.,  $k = \Omega(\frac{n}{(\alpha + \beta) \log(n)})$ , and as in the previous constructions from the previous sections, the size of the cut is  $O(\log(n))$ . Thus, the number of rounds needed for any algorithm to decide whether  $H$  is an  $(\alpha, \beta)$ -spanner of  $G$  is  $\Omega(\frac{n}{(\alpha + \beta) \cdot \log^3(n)})$ .

**Weighted Graphs** In order to achieve a higher lower bound for *weighted* graphs, we replace all the  $P$  and  $\frac{P}{2}$  paths by edges of weights  $P$  and  $\frac{P}{2}$  respectively. Note that in this case  $k = \Omega(n)$  and we deduce that the number of rounds needed for any algorithm to decide whether  $H$  is an  $(\alpha, \beta)$ -spanner of  $G$  is  $\Omega(\frac{n}{\log^2(n)})$ .

## 7 Discussion

We introduce a new technique for reducing the Set-Disjointness communication problem to distributed computation problems, in a highly efficient way. Our reductions encode an instance of Disjointness on  $k$  bits into a graph on only  $\tilde{O}(k)$  nodes and edges with a small “communication-cut” of size  $O(\log k)$ . All previous lower bound constructions had a cut of *poly*( $k$ ) size (e.g., [17, 21, 25, 26]). This efficiency allows us to answer several central open questions regarding the round complexity of important distance computation problems in the *CONGEST* model.

There are several interesting directions for future work. First, there is still a  $\log n$  factor gap between the upper and lower bounds on the round complexity of computing the diameter in the *CONGEST* model. Due to the fundamentality of the diameter, we believe that it will be interesting to close this small gap.

Second, we believe our degree-reduction technique can be used also for our bounds on approximations, albeit with more involved modifications. We also believe this technique can be useful in obtaining lower bounds on constant degree graphs in additional settings, beyond *CONGEST*.

Third, while our ideas greatly improve the state of the art lower bounds for shortest paths problems on *unweighted* graphs, their potential in the regime of *weighted* graphs is yet to be explored.

Finally, following our strong barriers for sparse graphs, it is important to seek further natural restrictions on the networks that would allow for much faster distance computation. Planar graphs are an intriguing setting in this context. A promising recent work of Ghaffari and Haeupler [22] showed that computing a minimum spanning tree can be done in  $\tilde{O}(D)$  rounds in planar graphs, despite the  $\tilde{\Omega}(\sqrt{n})$  lower bound for general graphs [17]. Can the diameter of a planar network be computed in  $\tilde{O}(D)$  rounds? While the graphs in our lower bounds are highly non-planar, it is interesting to note that they have a relatively small treewidth of  $O(\log n)$ .

**Acknowledgement:** We thank Ami Paz for pointing out Observation 6.4, and for useful discussions.

## References

- [1] A. Abboud, A. Backurs, T. D. Hansen, V. V. Williams, and O. Zamir. Subtree isomorphism revisited. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1256–1271, 2016.
- [2] A. Abboud, F. Grandoni, and V. V. Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1681–1697, 2015.
- [3] A. Abboud and V. V. Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 434–443, 2014.
- [4] A. Abboud, V. V. Williams, and J. R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 377–391, 2016.
- [5] A. Abboud, V. V. Williams, and O. Weimann. Consequences of faster alignment of sequences. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pages 39–51, 2014.
- [6] A. Abboud, V. V. Williams, and H. Yu. Matching triangles and basing hardness on an extremely popular conjecture. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 41–50, 2015.
- [7] D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.
- [8] A. Backurs and P. Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 51–58, 2015.
- [9] M. Borassi, P. Crescenzi, and M. Habib. Into the square - on the complexity of quadratic-time solvable problems. *CoRR*, abs/1407.4972, 2014.
- [10] M. Cairo, R. Grossi, and R. Rizzi. New bounds for approximating extremal distances in undirected graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 363–376, 2016.
- [11] M. L. Carmosino, J. Gao, R. Impagliazzo, I. Mihajlin, R. Paturi, and S. Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 261–270, 2016.

- [12] K. Censor-Hillel, P. Kaski, J. H. Korhonen, C. Lenzen, A. Paz, and J. Suomela. Algebraic methods in the congested clique. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*, pages 143–152, 2015.
- [13] K. Censor-Hillel, T. Kavitha, A. Paz, and A. Yehudayoff. Distributed construction of purely additive spanners. *Manuscript*, 2016.
- [14] T. M. Chan and R. Williams. Deterministic apsp, orthogonal vectors, and more: Quickly derandomizing razborov-smolensky. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1246–1255, 2016.
- [15] S. Chechik, D. H. Larkin, L. Roditty, G. Schoenebeck, R. E. Tarjan, and V. V. Williams. Better approximation algorithms for the graph diameter. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1041–1052, 2014.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.
- [17] A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012.
- [18] A. Drucker, F. Kuhn, and R. Oshman. On the power of the congested clique model. In *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 367–376, 2014.
- [19] M. Elkin. Unconditional lower bounds on the time-approximation tradeoffs for the distributed minimum spanning tree problem. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 331–340, 2004.
- [20] M. Elkin. An unconditional lower bound on the time-approximation trade-off for the distributed minimum spanning tree problem. *SIAM J. Comput.*, 36(2):433–456, 2006.
- [21] S. Frischknecht, S. Holzer, and R. Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1150–1162, 2012.
- [22] M. Ghaffari and B. Haeupler. Distributed algorithms for planar networks II: low-congestion shortcuts, mst, and min-cut. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 202–219, 2016.
- [23] M. Henzinger, S. Krinninger, and D. Nanongkai. An almost-tight distributed algorithm for computing single-source shortest paths. *CoRR*, abs/1504.07056, 2015.
- [24] S. Holzer, D. Peleg, L. Roditty, and R. Wattenhofer. Distributed  $3/2$ -approximation of the diameter. In *Distributed Computing - 28th International Symposium, DISC 2014, Austin, TX, USA, October 12-15, 2014. Proceedings*, pages 562–564, 2014.

- [25] S. Holzer and N. Pinsker. Approximation of distances and shortest paths in the broadcast congest clique. *CoRR*, abs/1412.3445, 2014.
- [26] S. Holzer and R. Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *ACM Symposium on Principles of Distributed Computing, PODC '12, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 355–364, 2012.
- [27] L. Kor, A. Korman, and D. Peleg. Tight bounds for distributed minimum-weight spanning tree verification. *Theory Comput. Syst.*, 53(2):318–340, 2013.
- [28] E. Kushilevitz and N. Nisan. Communication complexity. In *Cambridge University Press*, 1997.
- [29] C. Lenzen and B. Patt-Shamir. Fast routing table construction using small messages: extended abstract. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 381–390, 2013.
- [30] C. Lenzen and D. Peleg. Efficient distributed source detection with limited bandwidth. In *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 375–382, 2013.
- [31] J. Leskovec and A. Krevl. {SNAP Datasets}:{Stanford} large network dataset collection. 2014.
- [32] R. Meusel, S. Vigna, O. Lehmberg, and C. Bizer. The graph structure in the web - analyzed on different aggregation levels. *J. Web Science*, 1(1):33–47, 2015.
- [33] D. Nanongkai. Distributed approximation algorithms for weighted shortest paths. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 565–573, 2014.
- [34] D. Nanongkai, A. D. Sarma, and G. Pandurangan. A tight unconditional lower bound on distributed randomwalk computation. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 257–266, 2011.
- [35] D. Peleg. Distributed computing: a locality-sensitive approach. In *Society for Industrial Mathematics*, 2000.
- [36] D. Peleg, L. Roditty, and E. Tal. Distributed algorithms for network diameter and girth. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II*, pages 660–672, 2012.
- [37] D. Peleg and V. Rubinovich. A near-tight lower bound on the time complexity of distributed MST construction. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*, pages 253–261, 1999.
- [38] A. A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.
- [39] L. Roditty and V. V. Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 515–524, 2013.

- [40] R. Williams. Faster all-pairs shortest paths via circuit complexity. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 664–673, 2014.
- [41] V. V. Williams and R. Williams. Subcubic equivalences between path, matrix and triangle problems. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 645–654, 2010.
- [42] A. C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the 11h Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 209–213, 1979.